arXiv:1704.03296v1 [cs.CV] 11 Apr 2017

Interpretable Explanations of Black Boxes by Meaningful Perturbation

Ruth C. Fong
University of Oxford
ruthfong@robots.ox.ac.uk

Andrea Vedaldi
University of Oxford
vedaldi@robots.ox.ac.uk

Abstract

As machine learning algorithms are increasingly applied to high impact yet high risk tasks, e.g. problems in health, it is critical that researchers can explain how such algorithms arrived at their predictions. In recent years, a number of image saliency methods have been developed to summarize where highly complex neural networks "look" in an image for evidence for their predictions. However, these techniques are limited by their heuristic nature and architectural constraints.

In this paper, we make two main contributions: First, we propose a general framework for learning different kinds of explanations for any black box algorithm. Second, we introduce a paradigm that learns the minimally salient part of an image by directly editing it and learning from the corresponding changes to its output. Unlike previous works, our method is model-agnostic and testable because it is grounded in replicable image perturbations.

1. Introduction

Given the powerful but often opaque nature of modern black box predictors such as deep neural networks [4, 5], there is a considerable interest in *explaining and understanding* predictors *a-posteriori*, after they have been learned. This remains largely an open problem. One reason is that we lack a formal understanding of what it means to explain a classifier. Most of the existing approaches, such as DeConvNet [18, 15, 7], natural pre-images [6, 8], occlusion [18], etc., often produce intuitive visualizations; however, since such visualizations are primarily heuristic, their meaning remains unclear.

In this paper, we revisit the concept of "explanation" at a formal level, with the goal of developing principles and methods to explain any black box function f, e.g. a neural network object classifier. Since such a function is learned automatically from data, we would like to understand *what* f has learned to do and *how* it does it. Answering the "what" question means determining the properties of the map. The "how" question investigates the internal mechanisms that allow the map to achieve these properties. We focus mainly on the "what" question and argue that it can be answered by providing *interpretable rules* that describe the input-output relationship captured by f. For example, one rule could be that f is rotation invariant, in the sense that "$f(x) = f(x')$ whenever images x and x' are related by a rotation".

Figure 1. An example of a mask learned (right) by blurring an image (middle) to suppress the softmax probability of its target class (left: original image; softmax scores above images).

Such explanatory rules make *predictions* on the outcome of the black box f. Finding explanations can therefore be formulated as the problem of learning *meta-predictors*. In this paper, we investigate this concretely by focusing on the case of *local explanations* and reformulating *neural network saliency maps* in this framework (fig. 1). A local explanation tries to clarify *why* a black box makes a certain decision, say for example that image x_0 contains a *robin* ($f(x_0) = +1$). Since f is a black box, we must answer this question by observing its behaviour to *meaningful changes in the input*. For example, we can consider images x obtained from x_0 by *deleting* a subregion R from it. Depending on whether $f(x) = f(x_0)$, we can find out if R contains useful information for f. By finding all such regions, we can plot a saliency map of the image that identifies "what" is relevant to the black box.

In this paper, we make **several contributions**. First, we propose the general framework of **explanations as meta-predictors**, which extends previous related work ([17], section 2). Second, we identify several pitfalls in designing automatic explanation systems. We show in particular that **neural network artifacts** are a major attractor for explanations. While artifacts are informative since they explain part of the network behavior, characterizing other proper-

ties of the network requires careful calibration of the *generality* and *interpretability* of explanations. Third, we reinterpret network saliency in our framework. We show that this provides a natural **generalization of the gradient-based saliency technique** of [14] by *integrating* information over several rounds of backpropagation in order to *learn an explanation*. We also compare this technique to other methods such as guided backprop [15, 7] and excitation backprop [19] in terms of their meaning and obtained results.

2. Related work

Our work builds off of the gradient-based method introduced by [14], which backpropagates the gradient for a class label to the image layer. Other backpropagation methods include DeConvNet [18] and Guided Backprop [15, 7], which builds off of DeConvNet [18] and [14]'s gradient method to produce sharper visualizations.

Another set of techniques incorporate network activations into their visualizations: Class Activation Mapping (CAM) [21] and its relaxed generalization Grad-CAM [13] visualize the linear combination of a late layer's activations and class-specific weights (or gradients for [21]), while Layer-Wise Relevance Propagation (LRP) [1] and Excitation Backprop [19] backpropagate an class-specific error signal though a network while multiplying it with each convolutional layer's activations.

With the exception of [14]'s gradient method, the above techniques all arbitrarily invoke different backpropagation and/or activation-incorporating rules, which results in aesthetically pleasing, heuristic explanations of image saliency. They also are not model-agnostic, with most being limited to neural networks (all except [14, 1]) and many requiring specific architectural modifications [18, 15, 7, 21] and/or access to intermediate layers [21, 13, 1, 19].

A few techniques examine the relationship between inputs and outputs by editing an input image and observing its effect on the output. Two such methods include greedily graying out segments of an image until it is misclassified [20] and visualizing the drop in classification score when an image is occluded in various regions [18]. However, these techniques are limited by their approximate nature; we introduce a differentiable method that crucially allows for the effect of the joint inclusion or exclusion of different image regions to be considered.

In addition to [20, 18, 14], our research also builds off of work introduced in [17, 11, 2]. The idea of explanations as predictors is inspired by the work of [17]. However, we significantly generalize the latter to numerous new types of explanations, from classification to invariance.

The Local Intepretable Model-Agnostic Explanation (LIME) framework [11] is relevant to our local explanation paradigm and saliency method (sections 3.2 and 4) in that both use an algorithm's output with respect to inputs from a neighborhood around a given input x_0 that are generated by perturbing the image. However, their method takes much longer to converge ($N = 5000$ iterations compared to our 300 iterations) because they do not directly learn an image perturbation mask but rather sample perturbed inputs and learn the weights of a sparse linear classifier. Furthermore, their method limits learning to the top $K = 10$ superpixels per iteration. Additionally, we provide insight into the potential weaknesses of such methods; in contrast, [11] fails to consider how its method might learn artifacts from the hard contrast between included and excluded superpixels. Finally, their method does not yield a continuous heatmap that demonstrates the relative importance of different salient regions, which can be helpful for learning global patterns about class-specific features (section 5.4).

Similar to how our paradigm aims to learn an image perturbation mask that minimizes a class score, feedback networks [2] learn gating masks after every ReLU in a network to maximize a class score. However, our masks are plainly interpretable as they directly edit the image while [2]'s gates are not and can not be directly used as a visual explanation; furthermore, their method requires architectural modification and may yield different results for different networks, while ours is model-agnostic.

In short, we introduce a general framework for algorithmic explanations as well as a principled, fully-differentiable paradigm for learning a meaningful saliency map that explains which parts of an image most affect the output.

3. Explaining black boxes with meta-learning

A *black box* is a map $f : \mathcal{X} \to \mathcal{Y}$ from an input space $\mathcal{X}$ to an output space $\mathcal{Y}$, typically resulting from an opaque learning process. To make the discussion more concrete, consider as input color images $x : \Lambda \to \mathbb{R}^3$ where $\Lambda = \{1, \dots, H\} \times \{1, \dots, W\}$ is a discrete domain. The output $y \in \mathcal{Y}$ can be a boolean $\{-1, +1\}$ telling whether the image contains an object of a certain type (e.g. a *robin*), the probability of such an event, or some other interpretation of the image content.

In section 3.1, we develop a formal framework of explanations as *learnable meta-predictors*. In section 3.2, we use it to interpret [14]'s the gradient saliency technique as a *local explanation*. Then, in section 4, we develop this further into a new formulation of neural network saliency.

3.1. Explanations as meta-predictors

An *explanation* is a rule that predicts the response of a black box f to certain inputs. For example, we can explain a behavior of a *robin* classifier by the rule

$$Q_1(x; f) = \{x \in \mathcal{X}_c \Leftrightarrow f(x) = +1\} \tag{1}$$

where $\mathcal{X}_c \subset \mathcal{X}$ is the subset of all the *robin* images. Since f is imperfect, any such rule applies only approximately.

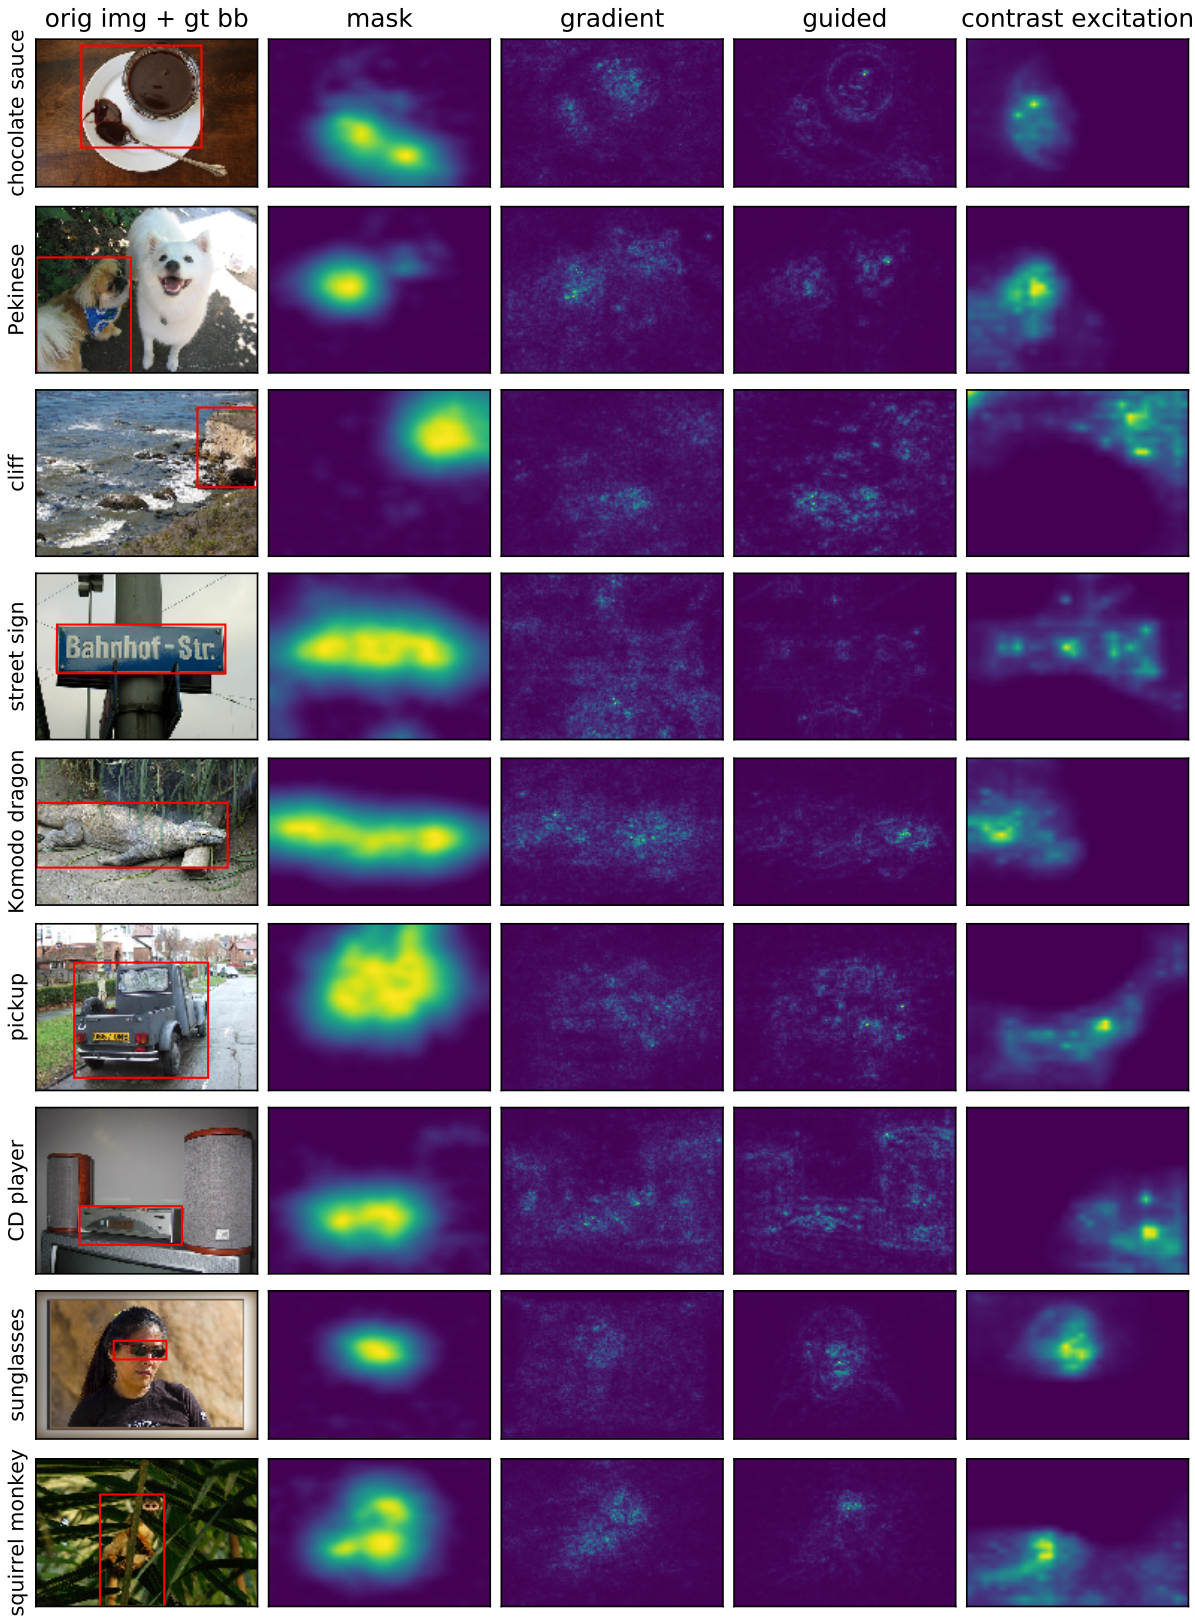

Figure 2. Comparison with other saliency methods. From left to right: original image with ground truth bounding box; learned mask subtracted from 1 (our method); gradient-based saliency [14]; guided backprop [15, 7]; contrastive excitation backprop [19].

We can measure the faithfulness of the explanation as its expected prediction error: $\mathcal{L}_1 = \mathbb{E}[1 - \delta_{Q_1(x;f)}]$ where δ_Q is the indicator function of event Q. Note that definition eq. (1) implicitly requires a distribution $p(x)$ over possible images $\mathcal{X}$. Note also that $\mathcal{L}_1$ is simply the expected prediction error of the classifier. Unless we did not know that f was trained as a *robin* classifier, explanation (1) is not very insightful, but it is interpretable since $\mathcal{X}_c$ is.

Explanations can also make relative statements about black box outcomes. For example, a black box f, could be rotation invariant:

$$Q_2(x, x'; f) = \{x \sim_{\text{rot}} x' \Rightarrow f(x) = f(x')\}, \quad (2)$$

where $x \sim_{\text{rot}} x'$ means that x and x' are related by a rotation. Just like before, we can measure the faithfulness of this explanation as $\mathcal{L}_2 = \mathbb{E}[1 - \delta_{Q_2(x,x';f)} | x \sim x']$.[1] This rule is interpretable because $\sim_{\text{rot}}$ has an obvious meaning.

Learning explanations. A significant advantage of formulating explanations as meta predictors is that their faithfulness can be measured as prediction accuracy. Furthermore, machine learning algorithms can be used to *discover explanations* automatically, by finding explanatory rules Q that apply to a certain classifier f out of a large pool of possible rules $\mathcal{Q}$.

In particular, finding the *most accurate* explanation Q is similar to a traditional learning problem and can be formulated computationally as a *regularized empirical risk minimization* such as:

$$\min_{Q \in \mathcal{Q}} \lambda \mathcal{R}(Q) + \frac{1}{n} \sum_{i=1}^{n} \mathcal{L}(Q(x_i; f), x_i, f), \; x_i \sim p(x). \quad (3)$$

Here, the regularization term $\mathcal{R}(Q)$ has two goals: to allow the explanation Q to generalize beyond the n samples $x_1, \ldots, x_n$ considered in the optimization and to pick an explanation Q which is simple and thus, hopefully, more interpretable.

Maximally informative explanations. Simplicity and interpretability are often not sufficient to find good explanations and must be paired with informativeness. Consider for example the following variant of rule Q_2:

$$Q_3(x, x'; f, \theta) = \{x \sim_\theta x' \Rightarrow f(x) = f(x')\}$$

where $x \sim_\theta x'$ means that x and x' are related by a rotation of up to an angle θ. Explanations for larger angles dominate (imply) the ones for smaller ones, with $\theta = 0$ being trivially satisfied. The regularizer $\mathcal{R}(Q_3(\cdot; \theta)) = -\theta$ can then be used to select a maximal angle and therefore find an explanation that is as informative as possible.[2]

[1] For rotation invariance we condition on $x \sim x'$ because the probability of independently sampling rotated x and x' is zero, so that, without conditioning, Q_2 would be true with probability 1.

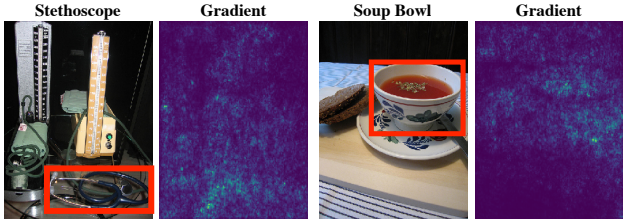

Figure 3. Gradient saliency maps of [14]. A red bounding box highlight the object which is meant to be recognized in the image. Note the strong response in apparently non-relevant image regions.

3.2. Local explanations

A *local explanation* is a rule $Q(x; f, x_0)$ that predicts the response of f in a neighborhood of a certain point x_0. If f is smooth at x_0, it is natural to construct Q by using the first-order Taylor expansion of f:

$$f(x) \approx Q(x; f, x_0) = f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle. \quad (4)$$

This formulation provides an interpretation of the *saliency maps* of [14], which proposed to visualize the gradient $S_1(x_0) = \nabla f(x_0)$ of the network an indication of salient image regions. They argue that large values of the gradient identify pixels that strongly affect the network output. However, we note here a limitation of this notion: *it breaks for a linear classifier*. In fact, if $f(x) = \langle w, x \rangle + b$, the saliency map $S_1(x_0) = \nabla f(x_0) = w$ is independent of the input image x_0.

The reason for this failure is that eq. (4) studies the variation of f for arbitrary displacements $\Delta_x = x - x_0$ from x_0 and, for a linear classifier, the change is the same regardless of the starting point x_0. For a non-linear black box f such as a neural network, this problem is reduced but not eliminated, and can explain why the saliency map S_1 is rather diffuse, with strong responses even where no obvious information can be found in the image (fig. 3).

We argue that this problem is a result of the fact that the meaning of explanations depends in large part on the *meaning of varying the input x to the black box*. For example, explanations in section 3.1 are based on letting x vary in an image category, or in pairs of rotated images. For saliency, one is interested in finding out regions in the image that are important for the network. Hence, it is natural to consider perturbations x obtained by deleting subregions of x_0. If we model deletion by multiplying x_0 point-wise by a mask

[2] Naively, strict invariance for any $\theta > 0$ implies invariance to arbitrary rotations as small rotations compose into larger ones. However, the formulation can still be used to characterize rotation insensitivity (when f varies slowly with rotation), or the meaning of $\sim_\theta$ can be changed to indicate rotation w.r.t. a canonical "upright" direction for a certain class of objects, etc.

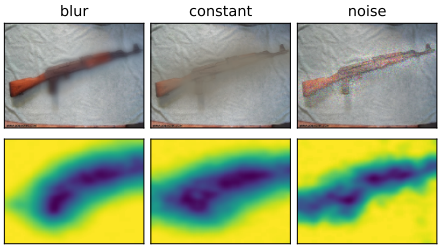

Figure 4. Perturbation types. Bottom: perturbation mask; top: effect of blur, constant, and noise perturbations.

m, this amounts to studying the function $f(x_0 \odot m)$[3]. The Taylor expansion of f at $m = (1, 1, \ldots, 1)$ is $S_2(x_0) = df(x_0 \odot m)/dm|_{m=(1,\ldots,1)} = \nabla f(x_0) \odot x_0$. For a linear classifier f, this results in the saliency $S_2(x_0) = w \odot x_0$ which is large for pixels for x_0 and w are large simultaneously. For a non-linear classifier, we refine this idea in the next section.

4. Saliency revisited

In this section we apply the general framework developed so far to the case, briefly considered in section 3.2, of generating saliency maps for black boxes.

4.1. Meaningful image perturbations

As argued above, in order to define an explanatory rule for a black box $f(x)$, one must start by specifying which variations of the input x will be used to study f. The aim of saliency is to identify which regions of an image x_0 are used by the black box to produce the output value $f(x_0)$. We can do so by observing how the value of $f(x)$ changes as x is obtained "deleting" different regions R of x_0. For example, if $f(x_0) = +1$ means that the image contains a robin, we expect that $f(x) = +1$ as well unless the choice of R deletes the robin from the image. Given that x is a perturbation of x_0, this is a local explanation (section 3.2) and we expect the explanation to characterize the relationship between f and x_0.

While conceptually simple, there are several problems with this idea The first one is to specify what it means "delete" information from an image region. As discussed in detail in section 4.3, we are generally interested in simulating naturalistic or plausible imaging effect, leading to more meaningful perturbations and hence explanations. Since we do not have access to the image generation process, we consider three obvious proxies: replacing the region R with a constant value, injecting noise, and blurring the image (fig. 4).

Formally, let $m : \Lambda \to [0, 1]$ be a *mask*, associating each pixel $u \in \Lambda$ with a scalar value $m(u)$. Then the perturbation operator is defined as

$$[\Phi(x_0; m)](u) = \begin{cases} m(u)x_0(u) + (1 - m(u))\mu_0, & \text{constant}, \\ m(u)x_0(u) + (1 - m(u))\eta(u), & \text{noise}, \\ \int g_{\sigma_0 m(u)}(v - u)x_0(v)\, dv, & \text{blur}, \end{cases}$$

where μ_0 is an average color, $\eta(u)$ are i.i.d. Gaussian noise samples for each pixel and σ_0 is the maximum isotropic standard deviation of the Gaussian blur kernel g_σ (we use $\sigma_0 = 10$, which yields a significantly blurred image).

4.2. Deletion and preservation

Given an image x_0, our goal is to summarize compactly the effect of deleting image regions in order to explain the behavior of the black box. A possible approach to this problem is to find deletion regions that are maximally informative.

In order to simplify the discussion, in the rest of the paper we consider black boxes $f(x) \in \mathbb{R}^C$ that generate a vector of scores for different hypotheses about the content of the image (e.g. as a softmax probability layer in a neural network). Then, we consider a "deletion game" where the goal is to find the smallest deletion mask m that causes the score $f_c(\Phi(x_0; m)) \ll f_c(x_0)$ to drop significantly, where c is the target class. Finding m can be formulated as the following learning problem:

$$m^* = \operatorname*{argmin}_{m \in [0,1]^\Omega} \lambda \|1 - m\|_1 + f(\Phi(x_0; m)) \qquad (5)$$

where λ is an energy term that encourages most of the mask to be turned off (hence deleting a small subset of x_0). In this manner, we can find a highly informative region for the network.

One can also play an symmetric "preservation game", where the goal is to find the smallest subset of the image that must be retained to preserve the score $f_c(\Phi(x_0; m)) \geq f_c(x_0)$: $m^* = \operatorname{argmin}_m \lambda \|m\|_1 - f(\Phi(x_0; m))$. The main difference is that the deletion game removes enough evidence to prevent the network from recognizing the object in the image, whereas the preservation game finds a minimal subset of sufficient evidence.

Iterated gradients. Both optimization problems are solved by using a local search by means of gradient descent methods. In this manner, our method extracts information from the black box f by computing its gradient, similar to the approach of [14]. However, it differs in that it extracts this information progressively, over several gradient evaluations, accumulating increasingly more information over time.

4.3. Dealing with artifacts

By committing to finding a single representative perturbation, deletion, preservation, and similar approaches incur

[3] $\odot$ is the Hadamard or element-wise product of vectors.

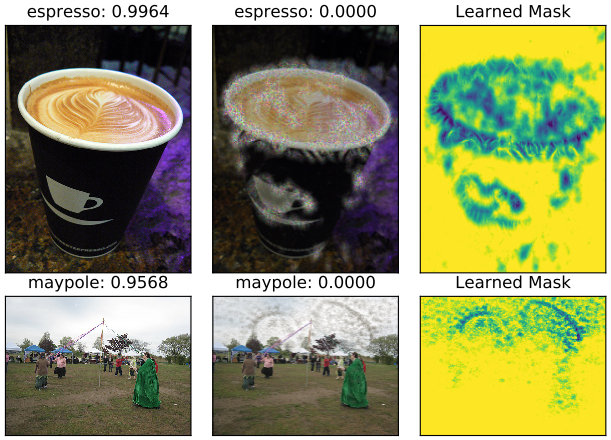

Figure 5. From left to right: an image correctly classified with large confidence by Googlenet [16]; a perturbed image that is not recognized correctly anymore; the deletion mask learned with artifacts. Top: A mask learned by minimizing the top five predicted classes by jointly applying the constant, random noise, and blur perturbations. Note that the mask learns to add highly structured swirls along the rim of the cup ($\gamma = 1, \lambda_1 = 10^{-5}, \lambda_2 = 10^{-3}, \beta = 3$). Bottom: A minimizing-top5 mask learned by applying a constant perturbation. Notice that the mask learns to introduce sharp, unnatural artifacts in the sky instead of "properly" deleting the pole ($\gamma = 0.1, \lambda_1 = 10^{-4}, \lambda_2 = 10^{-2}, \beta = 3$).

the risk of triggering artifacts of the black box. Neural networks, in particular, are known to be affected by surprising artifacts. For example, [5] show that a nearly-invisible adversarial perturbation of an image can lead a neural network to classify an object for another; [9] construct abstract synthetic images that are classified arbitrarily; and [6] find deconstructed versions of an image which are indistinguishable from the viewpoint of the neural network from the original image.

These examples demonstrate that it is possible to find particular inputs that can drive the neural network to generate nonsensical or unexpected outputs. This is not entirely surprising since neural networks are trained discriminatively on natural image statistics. While not all artifacts look "unnatural", nevertheless they form a subset of images that is sampled with negligible probability when the network is operated normally.

While the existence and characterization of artifacts is an interesting problem *per se*, we wish to characterize the behavior of black boxes under normal operating conditions. Unfortunately, as illustrated in fig. 5, objectives such as eq. (5) are strongly attracted by such artifacts, and naively learn subtly-structured deletion masks that trigger them. This is particularly true for the noise and constant perturbations as they can more easily than blur create artifacts using sharp color contrasts (fig. 5, bottom row).

We suggests two approaches to avoid such artifacts in generating explanations. The first one is that powerful explanations should, just like any predictor, generalize as much as possible. For the deletion game, this means not relying on the details of a singly-learned mask m. Hence, we reformulate the problem to apply the mask m stochastically, up to small random jitter.

Second, we argue mask co-adapted with network artifacts are *not representative of natural perturbations*. As noted before, the meaning of an explanation depends on the meaning of the changes applied to the input x; to obtain a mask more representative of natural perturbations we can encourage it to have a simple, regular structure which cannot be co-adapted to artifacts. We do so by regularizing m in total-variation (TV) norm and upsampling it from a low resolution version.

With these two modifications, eq. (5) becomes:

$$\min_{m \in [0,1]^\Lambda} \lambda_1 \|1 - m\|_1 + \lambda_2 \sum_{u \in \Lambda} \|\nabla m(u)\|_\beta^\beta + \mathbb{E}_\tau[f_c(\Phi(x_0(\cdot - \tau), M))], \quad (6)$$

where $M(v) = \sum_u g_{\sigma_m}(v/s - u)m(u)$. is the upsampled mask and g_{σ_m} is a 2D Gaussian kernel. Equation (6) can be optimized using stochastic gradient descent.

Implementation details. Unless otherwise specified, the visualizations shown were generated using Adam [3] to minimize Googlenet's [16] softmax probability of the target class by using the blur perturbation with the following parameters: learning rate $\gamma = 0.1, N = 300$ iterations, $\lambda_1 = 10^{-4}, \lambda_2 = 10^{-2}, \beta = 3$, upsampling a mask ($28 \times 28$ for Googlenet) by a factor of $\delta = 8$ and blurring the upsampled mask with $g_{\sigma_m=5}$. We initialize the mask as the smallest centered circular mask that suppresses the score of the original image by 99% when compared to that of the fully perturbed image, i.e. a fully blurred image.

5. Experiments

5.1. Interpretability

An advantage of the proposed framework is that the generated visualizations are clearly interpretable. For example, the deletion game produces a minimal mask that prevents the network from recognizing the object.

When compared to other techniques (fig. 2), this method can pinpoint the reason why a certain object is recognized without highlighting non-essential evidence. This can be noted in fig. 2 for the CD player (row 7) where other visualizations also emphasize the neighboring speakers, and similarly for the cliff (row 3), the street sign (row 4), and the sunglasses (row 8). Sometimes this shows that only a part of an object is essential: the face of the Pekenese dog (row 2), the upper half of the truck (row 6), and the spoon on the chocolate sauce plate (row 1) are all found to be minimally sufficient parts.

Figure 6. Interrogating suppressive effects. Left to right: original image with the learned mask overlaid; a boxed perturbation chosen out of interest (the truck's middle bounding box was chosen based on the contrastive excitation backprop heatmap from fig. 2, row 6); another boxed perturbation based on the learned mask (target softmax probabilities of for the original and perturbed images are listed above).

While contrastive excitation backprop generated heatmaps that were most similar to our masks, the fact that our paradigm is rooted in an interpretable edit to the image let's us demonstrate that our learned saliency maps are quantitatively better. In fig. 6, row 2, we show that applying a bounded perturbation informed by our learned mask significantly suppresses the truck softmax score, whereas a the boxed perturbation on the truck's back bumper, which is highlighted by contrastive excitation backprop in fig. 2, row 6, actually increases the score from 0.717 to 0.850.

The principled interpretability of our method also allows us to identify instances when an algorithm may have learned the wrong association. In the case of the chocolate sauce in fig. 6, row 1, it is surprising that the spoon is highlighted by our learned mask, as one might expect the sauce-filled jar to be more salient. However, manually perturbing the image reveals that indeed the spoon is more suppressive than the jar. One explanation is that the ImageNet "chocolate sauce" images contain more spoons than jars, which appears to be true upon examining some images. More generally, our method allows us to diagnose highly-predictive yet non-intuitive and possibly misleading correlations by identified machine learning algorithms in the data.

5.2. Deletion region representativeness

To test that our learned masks are generalizable and robust against artifacts, we simplify our masks by further blurring them and then slicing them into binary masks by thresholding the smoothed masks by $\alpha \in [0 : 0.05 : 0.95]$ (fig. 7; $\alpha \in [0.2, 0.6]$ tends to cover the salient part identified by the learned mask). We then use these simplified masks to edit a set of 5,000 ImageNet images with constant, noise, and blur perturbations. Using Googlenet [16],

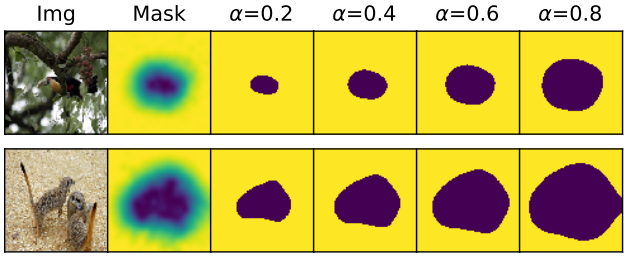

Figure 7. Left to right: original image, learned mask, and simplified masks for section 5.2 (not shown: further smoothed mask).

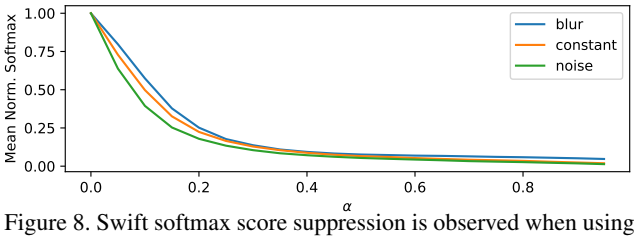

Figure 8. Swift softmax score suppression is observed when using all three perturbations with simplified binary masks (fig. 7) derived from our learned masks, thereby showing the generality of our masks.

we compute normalized softmax probabilities[4] (fig. 8). The fact that these simplified masks quickly suppress scores as α increases for all three perturbations gives confidence that the learned masks are identifying the right regions to perturb and are generalizable to a set of extracted masks and other perturbations that they were not trained on.

5.3. Minimality of deletions

In this experiments we assess the ability of our method to correctly identify a minimal region that suppresses the object. Given the output saliency map, we normalize its intensities to lie in the range $[0, 1]$, threshold it with $h \in [0 : 0.1 : 1]$, and fit the tightest bounding box around the resulting heatmap. We then blur the image in the box and compute the normalized[4] target softmax probability from Googlenet [16] of the partially blurred image.

From these sets of bounding boxes and normalized scores, for a given amount of score suppression, we find the smallest bounding box that achieves that amount of suppression. Figure 9 shows that, on average, our method yields the smallest minimal bounding box size when considering bounding boxes that suppress the normalized original score by $80\%, 90\%, 95\%,$ and 99%. These results show that our method finds a relatively minimal salient area that most impacts the output of the network.

5.4. Testing hypotheses: animal part saliency

From qualitatively examining learned masks for different animal images, we noticed that faces appeared to be

[4] $p' = \frac{p - p_0}{p_0 - p_b}$, where p, p_0, p_b are the masked, original, and fully blurred images' scores

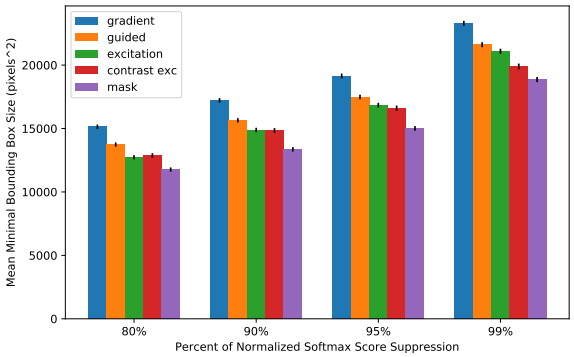

Figure 9. On average, our method generates the smallest bounding boxes that, when used to blur the original images, highly suppress their normalized softmax probabilities (standard error included).

Figure 10. "tiger" (top) and "bison" (bottom) images with eyes and feet annotations from [10]; our learned masks are overlaid. The mean average feet:eyes intensity ratio for "tigers" ($N = 25$) is 3.82, while that for bisons ($N = 22$) is 1.07.

more salient than appendages like feet. Because we produce dense heatmaps, we can test this hypothesis. From an annotated subset of the ImageNet dataset that identifies the keypoint locations of non-occluded eyes and feet of vertebrate animals [10], we select images from classes that have at least 10 images which each contain at least one eye and foot annotation, resulting in a set of 3558 images from 76 animal classes (fig. 10). For every keypoint, we calculate the average heatmap intensity of a 5×5 window around the keypoint. For all 76 classes, the mean average intensity of eyes were lower and thus more salient than that of feet (see supplementary materials for class-specific results).

5.5. Weakly Supervised Object Localization

Saliency methods are often assessed in terms of weakly-supervised localization. Because the deletion game is meant to discover minimal salient part and/or spurious correlation, we do not expect it to be particularly competitive on this task, but it is tested for completeness.

Similar to [19, 2], we predict a bounding box for the most dominant object in each of ∼50k ImageNet [12] validation images and employed three simple thresholding methods with which to fit a bounding box. First, for value thresholding, we normalize heatmaps to be in the range of $[0, 1]$ and then threshold them by their value with $\alpha \in [0 : 0.05 : 0.95]$. Second, for energy thresholding [2], we threshold heatmaps by the percentage of energy their most salient subset covered with $\alpha \in [0 : 0.05 : 0.95]$. Finally, with mean thresholding [19], we threshold a heatmap by $\tau = \alpha\mu_I$, where μ_I is the mean intensity of the heatmap and $\alpha \in [0 : 0.5 : 10]$. For each thresholding method, we search for the optimal α value on a heldout set. Localization error was calculated as the IOU with a threshold of 0.5.

	Val-α*	Err (%)	Ene-α*	Err	Mea-α*	Err
Grad [14]	0.25	46.0	0.10	43.9	5.0	41.7§
Guid [15, 7]	0.05	50.2	0.30	47.0	4.5	42.0§
Exc [19]	0.15	46.1	0.60	38.7	1.5	39.0§
C Exc [19]	—	—	—	—	0.0	57.0†
Feed [2]	—	—	0.95	38.8†	—	—
LRP [1]	—	—	—	—	1.0	57.8†
CAM [21]	—	—	—	—	1.0	48.1†
Mask‡	0.10	44.0	0.95	43.1	0.5	43.2

Table 1. Optimal α thresholds and error rates from the weak localization task on the ImageNet validation set using saliency heatmaps to generate bounding boxes. †Feedback error rate are taken from [2]; all others (contrastive excitation BP, LRP, and CAM) are taken from [19]. §Using [19]'s code, we recalculated these errors, which are $\leq 0.4\%$ of the originally reported rates. ‡Minimized top5 predicted classes' softmax scores and used $\lambda_1 = 10^{-3}$ and $\beta = 2.0$ (examples in supplementary materials).

Table 1 confirms that our method performs reasonably and shows that the three thresholding techniques used affect each saliency method differently. Our technique performs the best when using value thresholding and beats gradient and guided backprop when using energy thresholding. It also beats LRP, CAM, and contrastive excitation backprop when using mean thresholding (recall from fig. 2 that the contrastive method is visually the most similar to ours).

6. Conclusions

We propose a comprehensive, formal framework for learning algorithmic explanations as meta-predictors. We also present a novel image saliency paradigm that learns *where* an algorithm *looks* by discovering which parts of an image most affect its output score when perturbed. Unlike other saliency techniques, our method is grounded in a real edit to the image, making it uniquely interpretable and testable. We demonstrate numerous applications of our method, such as validating testable hypotheses (i.e. salient animal parts experiment) and comparing the suppressive effect of different saliency techniques (fig. 6 and fig. 9). Finally, we also contribute new insights into the fragility of neural networks and their susceptibility to artifacts, which opens a future research opportunity to adaptively augment networks to make them robust against undesirable artifacts.

References

[1] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015. 2, 8

[2] C. Cao, X. Liu, Y. Yang, Y. Yu, J. Wang, Z. Wang, Y. Huang, L. Wang, C. Huang, W. Xu, et al. Look and think twice: Capturing top-down visual attention with feedback convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2956–2964, 2015. 2, 8

[3] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6

[4] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1

[5] A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016. 1, 6

[6] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5188–5196, 2015. 1, 6

[7] A. Mahendran and A. Vedaldi. Salient deconvolutional networks. In *European Conference on Computer Vision*, pages 120–135. Springer International Publishing, 2016. 1, 2, 3, 8

[8] A. Mahendran and A. Vedaldi. Visualizing deep convolutional neural networks using natural pre-images. *International Journal of Computer Vision*, 120(3):233–255, 2016. 1

[9] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 427–436, 2015. 6

[10] D. Novotny, D. Larlus, and A. Vedaldi. I have seen enough: Transferring parts across categories. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016. 8

[11] M. T. Ribeiro, S. Singh, and C. Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. ACM, 2016. 2

[12] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 8

[13] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra. Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *arXiv preprint arXiv:1610.02391*, 2016. 2

[14] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *Proc. ICLR*, 2014. 2, 3, 4, 5, 8

[15] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014. 1, 2, 3, 8

[16] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. 6, 7

[17] R. Turner. A model explanation system. In *Proc. NIPS Workshop on Black Box Learning and Inference*, 2015. 1, 2

[18] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014. 1, 2

[19] J. Zhang, Z. Lin, J. Brandt, X. Shen, and S. Sclaroff. Top-down neural attention by excitation backprop. In *European Conference on Computer Vision*, pages 543–559. Springer, 2016. 2, 3, 8

[20] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014. 2

[21] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016. 2, 8